

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR U.S. LETTERS PATENT

TITLE:

A SYSTEM FOR AND METHOD OF CONTROLLING A VLSI ENVIRONMENT

INVENTORS:

Christopher A. Poirier
631 Blue Mountain Drive
Fort Collins, CO 80526
Citizenship: U.S.A.

Samuel D. Naffziger
3749 Ashmount Drive
Fort Collins, CO 80525
Citizenship: U.S.A.

Christopher J. Bostak
2909 Waterstone Court
Fort Collins, CO 80525
Citizenship: U.S.A.

A SYSTEM FOR AND METHOD OF CONTROLLING A VLSI ENVIRONMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is related to U.S. patent application serial no. [[attorney docket no. 200208754-1 (P742US)]], entitled “METHOD AND SYSTEM FOR CALIBRATION OF A VOLTAGE CONTROLLED OSCILLATOR (VCO);” U.S. patent application serial no. [[attorney docket no. 200208755-1 (P743US)]], entitled “SYSTEM AND METHOD FOR MEASURING CURRENT;” and, U.S. patent application serial no. [[attorney docket no. 200208728-1 (P745US)]], entitled “A METHOD FOR MEASURING INTEGRATED CIRCUIT PROCESSOR POWER DEMAND AND ASSOCIATED SYSTEM,” filed concurrently herewith, the disclosures of which are hereby incorporated by reference herein in their entirety.

BACKGROUND

[0002] Integrated circuit microprocessors or CPUs are typically designed for worst-case conditions that may include parameters that are critical to the VLSI design, such as frequency, power, voltage, current, and temperature. Some integrated circuit and CPU designs assume a standard set of conditions that require guard-banding. In these designs, the allowable operating conditions for the CPU are set so that the CPU design limits cannot be reached. For example, although a processor is capable of operating at 130 Watts under normal operating conditions, it may be guard-banded and hence specified to operate at 100 Watts to prevent the processor from exceeding the design limit.

[0003] In some designs, processors monitor a particular error condition and operate so as to not exceed that parameter. For example, a temperature measurement circuit having a trip point is used to notify the processor of a thermal problem. Such thermal monitoring circuits typically monitor only a single location on the processor’s integrated circuit. As a result, unmonitored sections of the integrated circuit may be operating at temperatures exceeding the design limits or those sections may be operating at a temperature well below the design limit when a monitored section trips the thermal warning. This type of thermal monitoring is not efficient and does not allow the processor to operate at optimal conditions.

[0004] In other designs, the processor is characterized across all operating conditions to determine a worse-case power or frequency value. The processor is then limited or guard-

banded to this worst-case condition, which may occur only under rarely used conditions. This prevents the processor from using more efficient power values and frequencies during typical operations.

[0005] The prior art solutions using guard-banding or external monitoring circuits are incapable of controlling the VLSI environment of the processor. Prior art circuits for monitoring discrete variables do not communicate with each other and, therefore, do not provide for VLSI parameter optimization across multiple variables. Additionally, such discrete circuits offer limited recourse to correct typical CPU problems such as high operating temperatures or high system power. For example, a prior art solution may provide a thermal trip circuit that completely disables a processor if an excessively high temperature is reached. This solution would be incapable of providing graceful performance throttling under such conditions.

SUMMARY

[0006] One embodiment of the invention includes a system comprising an integrated circuit on a VLSI die, and an embedded micro-controller constructed on the VLSI die, the micro-controller adapted to monitor and control the VLSI environment to optimize the integrated circuit operation.

[0007] Another embodiment of the invention includes a method for monitoring and controlling an integrated circuit comprising providing an embedded micro-controller on a same VLSI die as the integrated circuit, and monitoring and controlling a VLSI environment of the integrated circuit with the embedded micro-controller.

[0008] Another embodiment of the invention includes a computer program product comprising a computer usable medium having computer readable program code embedded therein, the computer readable program code comprising code for controlling an embedded micro-controller constructed on a VLSI integrated circuit die with a processor, wherein the micro-controller monitors and controls a VLSI environment of the processor.

[0009] An additional embodiment of the invention includes a system for monitoring and controlling an integrated circuit comprising means for providing an embedded micro-controller on a same VLSI die as the integrated circuit, and means for monitoring and controlling a VLSI environment of the integrated circuit with the embedded micro-controller.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIGURE 1 illustrates an embedded micro-controller according to one embodiment of the invention;

[0011] FIGURE 2 is a flowchart illustrating an exemplary process used by an embedded, on-die micro-controller to monitor the VLSI environment of an integrated circuit; and

[0012] FIGURE 3 illustrates a method for monitoring and controlling an integrated circuit.

DETAILED DESCRIPTION

[0013] A microprocessor system includes an embedded micro-controller that is constructed directly on the same integrated circuit die as a large VLSI CPU. The micro-controller's purpose is to control the VLSI environment, including, but not limited to, the power, temperature, voltage, current, frequency, and cooling air supply. The embedded, on-die micro-controller may employ a system of sensors and actuators to process the VLSI environment information, to determine an optimal operating solution, and to control the VLSI environment to achieve that solution.

[0014] The micro-controller may perform the following functions to monitor and control the integrated circuit environment: control power consumption, monitor and limit on-die temperature, adjust frequency based on voltage, adjust power supply voltage, and monitor die current consumption. Using the embedded micro-controller on the VLSI CPU die, the system can take many parameters into account for the particular die running in the context of a particular system environment. The micro-controller in some embodiments optimizes the VLSI parameters to provide an environment that will allow the CPU to operate as close to its design parameters as possible.

[0015] For example, the micro-controller may monitor voltage and current and may use those parameters to compute the system power. The micro-controller may use the power computation to adjust the power supply voltage as part of a feedback control system to control system power levels. The micro-controller may also be used as a digital filter to insure feedback stability of the power control loop.

[0016] The micro-controller may monitor temperature and may adjust power to gracefully limit on-die temperature. Alternatively, the micro-controller may adjust frequency

based on die voltage and temperature to prevent over-temperature conditions. The micro-controller may adjust voltage to the level that is required to support a given frequency.

[0017] The embedded micro-controller may consider all of the above-listed factors and more simultaneously and may use VLSI optimization algorithms that are implemented in firmware. The micro-controller provides advantages such as minimizing guard-banding, real time control and adjustment of the VLSI environment, flexibility to change the algorithms by re-programming the micro-controller firmware either to correct bugs or to offer customized solutions using software, and the ability to optimize across many variables. The use of an on-die micro-controller may enable a large VLSI CPU to adapt to and to control its specific operating environment.

[0018] FIGURE 1 illustrates an exemplary embedded micro-controller according to one embodiment of the invention. System 100 is a simplified, high-level block diagram illustrating a VLSI die for a CPU. The CPU includes two core processors 101 and 102 that are constructed on the same die as micro-controller 103. Each of the cores may include an integer unit and a floating point unit. Temperature sensors may be located near each integer unit and floating point unit. In core 101, temperature sensor 106 monitors integer unit 104 and temperature sensor 107 monitors floating point unit 105. In core 102, temperature sensor 109 monitors integer unit 108 and temperature sensor 111 monitors floating point unit 110.

[0019] In a preferred embodiment, the temperature sensors may be diodes coupled to a current source. The diodes are preferably sensitive to temperature and the voltage drop across the diode may vary with temperature, such as -1.7 mV/°C. Micro-controller 103 measures the voltage drop across the diode and uses the voltage information to calculate the temperature of the CPU core. The micro-controller may use analog-digital converters in the ammeters 112 to measure voltage.

[0020] The use of four separate temperature sensors allows micro-controller 103 to simultaneously monitor temperatures in different parts of the CPU and to get a more accurate measurement of the operating temperatures. Compared to prior art CPU designs, a lower threshold or maximum temperature, on the order of 90° C, can be used in the system of FIGURE 1. Micro-controller 103 may respond to high temperatures (i.e. temperatures approaching 90° C in any of the four sensors) by reducing the CPU power. Micro-controller 103 reduces the power by commanding voltage regulator module 113 to drop the power supply voltage or current that is provided to CPU cores 101, 102. In turn the available power will also drop. Ammeters 112 can

be used to measure the current and power to the CPU. Micro-controller 103 may also reduce the CPU's operating frequency by reducing the CPU's clock frequency. Clock system 114 reduces the system clock frequency under command by micro-controller 103.

[0021] A high temperature detected in one core may indicate that that core has a heavier workload compared to the other core. For example, if the temperature at sensor 106 in core 101 is approaching the maximum value, while the other temperature sensors remain at a lower level, it may indicate that integer unit 104 of core 101 has a heavy workload and that integer unit 108 has a relatively light workload. Upon detecting this difference in temperatures, micro-controller 103 may instruct the CPU's operating system to redistribute the workload so that integer unit 108 picks up some of integer unit 104's work load, thereby lowering the operating temperature of core 101.

[0022] As illustrated by the proceeding example, the micro-controller optimizes the working conditions of system 100. Rather than shutting down a core under high temperature conditions, micro-controller 103 monitors temperature increases and gracefully lowers the core's performance level to keep the core temperature below the maximum limit. In a preferred embodiment, micro-controller 103 optimizes the environment of system 100 to maintain a designated power level, such as 100 W. However, temperature considerations may override the 100 Watts power goal. Accordingly, micro-controller 103 may reject settings that would allow the CPU cores to operate at 100 W, and that would cause an over-temperature condition, i.e. over 90° C in one or both of the cores.

[0023] Micro-controller 103 includes firmware 115, which may comprise algorithms for determining how to respond to various temperature, power, and other parameters. Firmware 115 may be updated or replaced, for example by patch firmware, to fix "bugs" in system 101 or to provide a custom environment for the CPU. For example, system 100 may be ordinarily operated to maintain 100 Watts power and 90° C max temperature. However, in some applications these conditions may be unsuitable, such as in a system such as a blade server with many CPUs. It may be difficult to cool the system if there are many heat-generating components, such as CPUs. A user may install updated or customized firmware 115 in micro-controller 103 so that, for example, system 100 is optimized to operate at a power level less than 100 Watts, such as 50 Watts, or at a maximum temperature less than 90° C.

[0024] In addition to software configuration information provided by firmware 115, fuses 116a-c provide hardware configuration control for micro-controller 103. If micro-

controller 103 senses a voltage across one or more of fuses 116a-c, then micro-controller 103 will (or will not) provide optimization control for that parameter. For example, if temperature fuse 116a is not blown and micro-controller senses a voltage on that line, then micro-controller 103 will provide temperature control to processor cores 101 and 102 in system 100. In an alternative embodiment, micro-controller 103 provides temperature control to system 100 if no voltage is sensed across temperature fuse 116a. Similarly, the voltages appearing across fuses 116b and 116c may impact whether micro-controller 103 provides power and voltage control to system 100. Other fuses (not shown) may provide a hardware configuration for micro-controller 103 to control other system parameters.

[0025] Micro-controller 104 uses ammeters 112, which may be high-precision voltmeters, to measure CPU power. Ammeters 112 are used to calculate the current flowing into the CPU by measuring the voltage drop across a parasitic resistance, such as the resistance of the CPU package or the resistance of the power supply grid. Alternatively, micro-controller 103 may use a predetermined resistance value or may calculate the parasitic resistance, for example, through a calibration operation. The voltage and resistance values are used to calculate current and power for the CPU. A method and system for calibrating ammeters on a CPU die is disclosed in concurrently filed, copending U.S. patent application serial no. [[attorney docket no. 200208728-1]], entitled A METHOD OF AND SYSTEM FOR CONTINUOUS ON-DIE AMMETER CALIBRATION TO COMPENSATE FOR TEMPERATURE AND DRIFT ON-BOARD A MICROPROCESSOR, the disclosure of which is hereby incorporated by reference herein.

[0026] Micro-controller 103 may control the clock frequency in system 100 by adjusting the available voltage from the power supply. The clock frequency provided by clock system 114 is proportional to the available system voltage. As micro-controller 103 reduces the voltage, the frequency of the clock signal is reduced by clock system 114 to a rate that can be supported by the available power. As a result, cores 101 and 102 perform fewer operations per second when the power is lowered, which causes the temperature of the processor cores to drop. In other embodiments, micro-controller can control the clock frequency directly without adjusting the system voltage or power. This may result in a less-than-optimum configuration regarding the relationship between the VLSI environment's power and frequency, but may be desired in certain instances.

[0027] Micro-controller 103 is capable of considering many parameters simultaneously and using those parameters to optimize the CPU operation. Micro-controller 103 may consider the power, voltage, current, temperature, and frequency parameters of the CPUs current operating condition. Using the optimization algorithms in firmware 115, micro-controller 103 adjusts selected parameters to ensure that system 100 does not go into an over-temperature condition while maintaining operations at or near the design system power level.

[0028] Although the system illustrated in FIGURE 1 shows two processor cores on a single die, those of skill in the art will understand that micro-controller 103 can also be used to control CPUs that comprise multiple dies and/or that include more than two processor cores on one or more dies. Moreover, micro-controller 103 can also be used to separately control the temperature and power for processor cores on multiple CPU dies. Furthermore, it will be understood by those of skill in the art that the present invention is not limited for monitoring and controlling processors or CPUs, but can be used to monitor and control the environment of any type of integrated circuit.

[0029] FIGURE 2 is a flowchart illustrating an exemplary process used by an embedded, on-die micro-controller to monitor and control the VLSI environment of a CPU. The micro-controller calibrates its sensors and/or look-up tables in 201. In a preferred embodiment, the calibration is part of an iterative process in which the micro-controller interleaves calibration steps with sensor measurements so that the calibration process does not interfere with the micro-controller's duties of monitoring and controlling the CPU environment.

[0030] In 202, the micro-controller monitors one or more temperature sensors for an over-temperature condition. If one or more temperature sensors indicate that an over-temperature condition exists, then process 200 moves to block 203 wherein the micro-controller reduces the clock frequency in an attempt to reduce the temperature of the processor core. In an alternative embodiment, at block 203, the micro-controller may reduce the CPU voltage which causes the clock frequency to decrease.

[0031] The micro-controller may detect an existing over-temperature condition at block 202. Alternatively, the micro-controller may compare a series of temperature readings to anticipate an over-temperature condition. For example, if each temperature measurement in a sequence of samples is higher than the previous measurement, then the micro-controller may react to prevent an expected maximum temperature. If the core temperature is within acceptable limits at block 202, but an over-temperature condition is projected, then micro-controller can

anticipate the temperature problem and will move to block 203 to reduce the system frequency in order to avoid the over-temperature condition.

[0032] If the temperature is acceptable in block 202, or after the clock frequency has been reduced in 203, then the micro-controller measures the frequency in block 204. If the clock frequency is below a desired range, then micro-controller increases the clock frequency in block 205. To prevent the micro-controller from counteracting a prior frequency correction, such as a correction in block 203, the micro-controller considers any current temperature measurement corrections that are in place prior to increasing the clock frequency in block 204.

[0033] If the frequency is within an acceptable range at block 204, or after the frequency is corrected in block 205, the micro-controller measures the CPU power level in block 206. If the power level is within an acceptable range, the process begins again at 201. If the power level is below an optimal range, then the micro-controller increases the power level in block 207 and repeats the process. If the CPU power level is above an optimal range at block 206, then the micro-controller decreases power in block 208 and repeats the process. The micro-controller attempts to maintain the CPU operating at its design power level. However, high temperature conditions detected in block 202 may prevent the micro-controller from increasing the power level.

[0034] After the power level is checked and adjusted as needed in blocks 206 and 207, the micro-controller returns to the calibration operation at 201 where it performs another calibration operation prior to commencing another pass through the CPU environment monitoring operations.

[0035] FIGURE 3 illustrates a method for monitoring and controlling an integrated circuit. The method comprises providing an embedded micro-controller on a same VLSI die as the integrated circuit, 301. The method further comprises monitoring and controlling a VLSI environment of the integrated circuit with the embedded micro-controller, 302.